

Modeling the Dynamics of Inflation Compensation*

Markus Jochmann

Department of Economics

University of Strathclyde

email: Markus.Jochmann@strath.ac.uk

Gary Koop

Department of Economics

University of Strathclyde

email: Gary.Koop@strath.ac.uk

Simon M. Potter

Macroeconomics and Monetary Studies Function

Federal Reserve Bank of New York

email: simon.potter@ny.frb.org

June 2008

Abstract

This paper investigates the relationship between short-term and long-term inflation expectations using daily data on inflation compensation. We use a flexible econometric model which allows us to uncover this relationship in a data-based manner. We relate our findings to the issue of whether inflation expectations are anchored, unmoored or contained. Our empirical results indicate no support for either unmoored or firmly anchored inflation expectations. Most evidence indicates that inflation expectations are contained.

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System. Markus Jochmann and Gary Koop are Fellows of the Rimini Center for Economic Analysis. Preliminary version, please do not quote.

1 Introduction

Many decisions of central bankers and models of economists depend on an understanding of how agents form expectations about inflation. In this paper, we shed light on this issue using data on inflation compensation derived from real and nominal U.S. Treasury securities. With some qualifications, one can interpret inflation compensation as a measure of expected inflation. Of particular interest is inflation pass through: how changes in short-term inflation expectations influence long-term expectations. We show how different models of inflation expectations, which we refer to as anchored, unmoored and contained, imply different forms for inflation pass through. We then investigate whether there is empirical support for any of these models.

Our empirical work uses a flexible parametric econometric framework developed in Koop and Potter (2007). This framework combines some of the benefits of a nonparametric approach (i.e. it lets the data speak rather than enforcing a parametric, often linear, functional form) with the benefits of a parametric model rooted in the considerations of economic theory. That is, different theoretical models of inflation expectations imply the inflation pass through coefficient should have various forms. We discuss how, if inflation expectations are anchored, the inflation pass through coefficient should be constant and small (and certainly less than one). Unmoored inflation expectations imply that the pass through coefficient should be near one. However, if inflation expectations are contained, then the magnitude of the pass through coefficient should vary with the level of short-term inflation expectations in a particular way. Such theoretical considerations can be incorporated in our flexible parametric framework and allow the data to tell us which (if any) specification is preferred.

Our empirical results, using daily data, indicate support for contained inflation expectations. We provide strong evidence that the inflation pass through coefficient is not constant, but does tend to be quite small. These findings are not consistent with either unmoored or anchored inflation expectations, but are consistent with the idea that inflation expectations are contained in some way. The particular theoretical model of contained expectations discussed in this paper suggests that the inflation pass through coefficient should depend on the absolute deviation of short term inflation expectations from a central value. Furthermore, it should have a certain functional form. We find empirical evidence in support of both these features of this theoretical model.

The paper is organized as follows. In the next section we define basic concepts and discuss the implications of various models of inflation expectations for the inflation pass through coefficient. In the third section, we describe and motivate our econometric methods (with technical details available in an appendix). The fourth section presents empirical results and the fifth section concludes.

2 Modeling Inflation Pass Through

In this paper we investigate the relationship between long-term expected inflation (y_t) and short-term expected inflation (x_t) using daily time series data ($t = 1, \dots, T$). We use inflation compensation as a measure of these expectations. We begin by defining these concepts and motivating why this relationship is of importance. Inflation compensation is calculated using interest rates on nominal and inflation-indexed bonds. Let $r_t^{(h)}$ is the real interest rate on an h -period bond (i.e. the real return that an investor requires to lend money at time t in return for a payment h years in the future) and $i_t^{(h)}$ be the nominal interest rate on an h -period bond. These variables reflect the total return between the present and m periods in the future. Using these returns, we can calculate inflation compensation between h_1 and h_2 periods in the future as:

$$ei^{(h_1, h_2)} = \frac{h_2 (i^{(h_2)} - r^{(h_2)}) - h_1 (i^{(h_1)} - r^{(h_1)})}{h_2 - h_1}$$

Inflation compensation is commonly used as a measure of expected inflation and we will refer to every one of the variables used in our empirical work as expected inflation for particular values of h_1 and h_2 . The issues under study involve the relationship between long-term and short-term expected inflation. Our empirical section defines long-term expected inflation as 9 to 10 years and short-term expected inflation as 2 to 5 years. Thus, we will use $y = ei^{(9,10)}$ and $x = ei^{(2,5)}$.

As a terminological digression, in this paper we informally use the terms expected inflation and inflation compensation (also known as breakeven inflation) interchangeably. Inflation compensation is the compensation that investors require for holding nominal rather than real inflation-indexed bonds. This compensation largely reflects expected inflation (and empirical studies such as ours typically interpret it as such). However, inflation compensation

also reflects compensation that investors are demanding for risks associated with the uncertainty about future inflation. But, following most of the literature (e.g. Gurkaynak, Levin and Swanson, 2006), we argue that the inflation expectation component will dominate (especially when using differenced data) and simply use the terminology “expected inflation” in this paper.

Important issues of policy hinge on inflation pass through (i.e. how changes in short term inflation expectations affect long term expectations¹). As a starting point, a simple regression model, of the sort which has been used in the literature (see, e.g., Potter and Rosenberg, 2007), can be used to investigate these issues:

$$\Delta y_t = \beta \Delta x_t + \varepsilon_t, \tag{1}$$

where β is referred to as the pass through coefficient. There are strong theoretical and empirical reasons for thinking that the pass through coefficient might not simply be a constant but might depend on the level of expected inflation or might be varying over time. These reasons motivate the flexible parametric treatment used in this paper which allows for the pass through coefficient to potentially have such properties. Before we describe our statistical model, it is useful to briefly describe some theoretical models of expected inflation (largely taken from Potter and Rosenberg, 2007) and discuss what their implications are for the pass through coefficients. This will allow us to formally define the concepts of anchored, contained and unmoored inflation expectations and help us later when it comes to interpreting our empirical results.

We begin with a standard decomposition of observed inflation (π_t) into permanent (π_t^*) and transitory (c_t) components:

$$\pi_t = \pi_t^* + c_t.$$

It is common to interpret π_t^* as underlying inflation and define it through the properties:

¹Note that, theoretically, short term inflation expectations should be very close to actual inflation. Given a lack of daily data on inflation, short term inflation expectations can be considered as a proxy for inflation. Thus, the inflation pass through coefficient can also be considered as an approximate measure of the effect of changes in current inflation on long term expectations of inflation.

$$\begin{aligned} E_t(\pi_{t+h}) &\rightarrow E_t(\pi_{t+h}^*) \\ E_t(c_{t+h}) &\rightarrow 0 \text{ as } h \rightarrow \infty. \end{aligned}$$

The properties of the pass through coefficient depend on underlying inflation. Potter and Rosenberg (2007) distinguish between anchored, contained and unmoored inflation expectations (and other papers adopt similar terminologies) and work out the implications for the pass through coefficient of several simple theoretical models of the inflation process. In the remainder of this section, we briefly summarize some of their results. In this theoretical discussion, we will let β_{h_1, h_2} be the pass through of changes in inflation expectations at horizon h_1 to changes in inflation expectations at horizon h_2 (where h_1 and h_2 are appropriately chosen to be short-term and long-term, respectively). We emphasize that, unlike a constant parameter model such as (1), we are not assuming β_{h_1, h_2} to necessarily be a single constant parameter. Its magnitude could vary with expected inflation or be time varying.

Anchored inflation expectations can arise when there is a credible inflation target ($\widehat{\pi}$). Faust and Henderson (2004) consider a stationary model of underlying inflation:

$$\pi_t^* = \widehat{\pi}(1 - \theta) + \theta\pi_{t-1}^* + u_t,$$

where u_t is a stationary residual and $|\theta| < 1$. It can be immediately seen that long run inflation expectations in this model will be $\widehat{\pi}$ and, thus, inflation expectations are anchored about this target. As shown in Potter and Rosenberg (2007), in this model the pass through coefficient is constant, does not depend on the level of expected inflation and is less than one in absolute value. To be precise, it is:

$$\beta_{h_1, h_2} = \theta^{h_2 - h_1}.$$

This illustrates how a constant pass through coefficient which is less than one in absolute value is consistent with anchored expectations. Furthermore, if h_1 is small (as with short-term inflation expectations) and h_2 is large (as will long term inflation expectations). then the pass through coefficient will tend to be small.

Unmoored inflation expectations arise through unobserved components models where underlying inflation is a random walk:

$$\pi_t^* = \pi_{t-1}^* + u_t.$$

This model, extended to allow for stochastic volatility, is used in Stock and Watson (2007). Potter and Rosenberg (2007) show that the model used in Stock and Watson (2007) implies:

$$\beta_{h_1, h_2} \rightarrow 1 \text{ as } h_2 \rightarrow \infty \text{ with } h_1 \text{ fixed.}$$

For finite h_2 the formula for β_{h_1, h_2} depends on the variances and persistences of the permanent (π_t^*) and transitory (c_t) components of inflation. However, a key finding is that, if h_2 is large enough so that the transitory shock has dissipated, then we have $\beta_{h_1, h_2} = 1$. Hence, we can take such models of unmoored inflation expectations as implying constant pass through coefficients of one (or close to one for finite h_2).

Potter (2007) proposes a third class of models which imply contained inflation expectations. This model has properties which lie between those of the anchored expectations and unmoored expectations models. Inflation expectations are not fully anchored, but are contained within bounds. In this model, underlying inflation expectations look like the unmoored expectations model:

$$\pi_t^* = \pi_{t-1}^* + u_t,$$

but u_t is constrained so that underlying inflation lies in a bounded interval, $[a, b]$. In particular,

$$u_t \sim TN(a - \pi_{t-1}^*, b - \pi_{t-1}^*; 0, \sigma_u^2),$$

where $TN(a - \pi_{t-1}^*, b - \pi_{t-1}^*; 0, \sigma_u^2)$ is the truncated Normal distribution (i.e. the $N(0, \sigma_u^2)$ distribution truncated to the interval $[a - \pi_{t-1}^*, b - \pi_{t-1}^*]$). Potter and Rosenberg (2007) show how this model implies that the pass through coefficient depends on short-term inflation expectations. Near either a or b , we have β_{h_1, h_2} being near zero. But, unlike models of anchored expectations, β_{h_1, h_2} becomes larger near the middle of the interval $[a, b]$. Empirically, we will take this theoretical model as implying that β_{h_1, h_2} is potentially large (i.e. it could even be near one) when short term inflation expectations are near a central value, $\bar{\pi}$ (e.g. the average value of inflation), but as short term inflation expectations deviate from $\bar{\pi}$, β_{h_1, h_2} goes to becomes smaller. Intuitively, agents believe that if short term inflation expectations are near

$\bar{\pi}$ the Fed will not act and underlying inflation could potentially just wander according to a random walk (and, thus, we obtain the $\beta_{h_1, h_2} \rightarrow 1$ result of the unmoored inflation expectations model). However, if short term inflation expectations deviate too far from $\bar{\pi}$ then the Fed will act and reversion will kick in (and we obtain a result closer to that obtained from the anchored expectations model where the pass through coefficient is small). If this model of contained inflation expectations is correct, then empirically we should have the magnitude of the pass through coefficient varying with the absolute deviation of short term inflation expectations from $\bar{\pi}$. In particular, the pass through coefficient could be large when $\left| ei_t^{(2,5)} - \bar{\pi} \right|$ is near zero, but should be small when it is large.

In our empirical work, we use the theory discussed in this section to guide our econometric specification. We adopt a statistical approach which lets the data speak, telling us what the form of the pass through coefficient is. We indirectly use this theory to motivate the idea the pass through coefficient could potentially depend on the level of short term inflation expectations. Furthermore, we use this theory to guide us in the interpretation of our results. That is, if the pass through coefficient is constant and small, that is consistent with anchored inflation expectations. If the pass through coefficient is always near one, that is consistent with unmoored inflation expectations. If the pass through coefficient is large at (approximately) average levels of short-term inflation expectations, but becomes small as short-term expectations deviate from average, then our findings are consistent with Potter (2007)'s model of contained inflation expectations. Our statistical methods, described in the next section, allow us to uncover which of these models of inflation expectations is most consistent with the data (as opposed to working with a model which imposes one particular view of inflation expectations). They can also tell us if the data is inconsistent with any of these models of inflation expectations.

3 Flexible Parametric Modeling

In previous work (see Koop and Potter, 2007), we developed a simple modeling framework that is extremely flexible, exactly or approximately nesting a wide variety of nonlinear time series and structural break models. We use this framework here. For present purposes, the key things to note are that it allows for the inflation pass through coefficient to be constant or change

over time (either in a gradual fashion or in a more abrupt fashion). It can also allow for the magnitude of the pass through coefficient to depend on another explanatory variable (which is referred to as an index variable). The arguments in the previous section suggest a logical choice for the index variable: short term inflation expectations (although we experiment with other choices as well). And the functional form describing the relationship between the index variable and the pass through coefficient can take on a wide range of shapes, depending on whether inflation expectations are anchored, unmoored or contained (or something else). Our flexible parametric approach seems ideally designed for the present problem. It allows the data to determine the relationship between long term and short term inflation expectations in an extremely flexible fashion. Loosely speaking, our approach can be thought of as “nonparametric in spirit”, but ideas drawn from economic theory (primarily through choice of index variables) are used to suggest empirically-useful directions for the model to look. In this section, we briefly describe our model. Bayesian methods for carrying out econometric inference in this model are described in the appendix. For complete details (and additional motivation for this flexible parametric approach), see Koop and Potter (2007).

As a first step, consider extending (1) to be a time varying parameter (TVP) regression model, written in state space form with measurement equation given by:

$$\Delta y_t = \beta_t \Delta x_t + \sigma_\varepsilon \varepsilon_t, \quad (2)$$

and state equation given by:

$$\beta_t = \beta_{t-1} + \sigma_v v_t \quad (3)$$

where ε_t and v_t are i.i.d. $N(0, 1)$ (and independent of one another). This model is of interest in and of itself (e.g. it allows for the pass through coefficient to gradually evolve over time, which might be of empirical relevance), nests some interesting models (e.g. $\sigma_v = 0$ is the linear model) and textbook methods for statistical inference are available.

What the flexible parametric approach of Koop and Potter (2007) does is add two extensions to this familiar framework. These extensions are both simple (and textbook methods of statistical inference require only very minor adaptations), but hugely increase the flexibility of the model. These extensions involve data re-ordering and distance between observations.

To understand what we mean by “data re-ordering”, note that the preceding state equation was written in natural time ordering ($t = 1, \dots, T$). But

we can also write a state equation with the observations ordered in a different manner:

$$\beta_s = \beta_{s-1} + \sigma_v d(z_s, z_{s-1}) v_s, \quad (4)$$

where s indexes a re-ordering (in ascending order) of time according to an *index variable* z_t . For instance, in our case z_t could be short term inflation expectations in which case, we re-order the data as $z_s = 1, \dots, T$ with z_1 having the lowest inflation expectations, z_2 the second lowest, etc. To aid in interpretation of the distance function, note that z_s and z_{s-1} will be adjacent observations (e.g. if the index is short term expected inflation, then z_{s-1} will have the next lowest value of short term expected inflation to z_s). $d(z_s, z_{s-1})$ is a (non-negative) distance function measuring the distance between adjacent observations, z_s and z_{s-1} .

These two apparently minor extensions add hugely to the flexibility of the model. To see the benefit of data re-ordering, consider a simple structural break model where the pass through coefficient changes at time τ (so that it is $\beta^{(1)}$ for $t = 1, \dots, \tau$ and is $\beta^{(2)}$ for $t = \tau + 1, \dots, T$). Next consider the case where the pass through coefficient takes on the value $\beta^{(1)}$ if short term inflation expectations are less than τ and is otherwise $\beta^{(2)}$. Our methods draw on the fact that these two cases are statistically exactly the same model, except that one uses the data in natural time ordering and the other would have all the variables ordered according to short term inflation expectations. This is an example of a general result that has long been recognized (see, e.g., Tsay, 1989): many common regime-switching and threshold regression models are equivalent to regression models with structural breaks if the data is suitably re-ordered. In words, the state equation for a conventional TVP regression model says “pass through coefficients *at similar points in time* are likely to be similar to one another”. Our state equation, (4), says “pass through coefficients *for observations with similar short term inflation expectations* are likely to be similar to one another” (when we use short-term inflation expectations as the index variable). In general, it implies “pass through coefficients for observations with similar values for the index variable are likely to be similar to one another”. Nonparametric regression models estimate the relationship $y = f(x)$ using the idea that observations with similar values for x imply similar values for $f(x)$ (e.g. kernel algorithms smooth over neighboring observations). In this sense, our flexible parametric methods are doing something similar and are nonparametric in spirit (except

that, unlike many nonparametric approaches, we do not choose a single index variable to define what we mean by “neighboring observations”).

Thus, simply by allowing for various re-orderings of the data, we can obtain specifications such as those in Potter and Rosenberg (2007). For instance, their model where inflation expectations are contained, implies the pass through coefficient depends on expected inflation (being highest when expected inflation is near a central value, but declining as it moves away from target). Such a pattern can easily be obtained using our model. But our model is much more flexible, allowing for different functional forms than the inverted-U of Figure 2 of Potter and Rosenberg (2007). Koop and Potter (2007) explore just how flexible this approach is. But, it is worth noting in passing that the relationship between state space models and nonparametric kernel smoothing algorithms is well-developed in the state space literature (e.g. Harvey, 1989 and Harvey and Koopman, 2000).

An important feature of our approach is that it does not involve selecting just one index variable. Instead, it allows for a variety of index variables to be selected. One can treat it as either a model selection device (i.e. the researcher can select the index variable which receives most support and use that model) or as a model averaging device (i.e. the researcher can obtain results using each index variable and then average over them with data-based weights²). From a purely statistical perspective, there is no reason not to consider large numbers of different index variables. Indeed, if we considered every possible way of re-ordering the data, this class of models would be so flexible as to be virtually equivalent to a nonparametric model. However, from an economic perspective, there are strong reasons to limit the set of index variables. Empirical considerations suggest that time variation in pass through is worth investigating.³ The theoretical discussion of the preceding section suggests that it is worth investigating whether the magnitude of the pass through coefficient is related to short term inflation expectations. Accordingly, we use the following choices for z :

1. $z_t = t$ is natural time ordering
2. $z_t = e i_t^{(2,5)}$

²In standard Bayesian fashion, these data based weights are proportional to marginal likelihoods.

³Alternatively, it might be the current or lagged level of inflation which should be the index variable. Given the lack of daily data on inflation, this can be proxied by short term inflation expectations.

3. $z_t = \left| ei_t^{(2,5)} - \bar{\pi} \right|$, where $\bar{\pi} = 2.25\%$
4. $z_t = ei_{t-1}^{(2,5)}$
5. $z_t = \left| ei_{t-1}^{(2,5)} - \bar{\pi} \right|$.

The second extension of the standard TVP model is to allow for a distance function, $d(z_s, z_{s-1})$, in (4). This allows for even greater flexibility. In particular, note that standard TVP state equations such as (3) (or our state equation, 4, with $d(z_s, z_{s-1}) = 1$ for all s) allow for continual gradual evolution of the pass through coefficient. By adding the distance function we relax the “continuous gradual” aspect. Our model does allow for continuous gradual change, but we can also have either irregular change in the pass through coefficient (e.g. if it is constant over long spells of time) or change can be more abrupt. To illustrate the latter point, suppose the index variable is a measure of short-term inflation expectations and we have $d(z_s, z_{s-1}) = 1$ if $z_{s-1} < \tau$ and $z_s \geq \tau$, and $d(z_s, z_{s-1}) = 0$ otherwise. Then we have a model where the pass through coefficient changes abruptly once short-term inflation expectations hit a threshold, τ .

As argued in Koop and Potter (2007), a flexible distance function is:

$$d(z_s, z_{s-1}) = \Phi(z_s) - \Phi(z_{s-1}),$$

where $\Phi(z)$ is the c.d.f. of the $N(\mu_d, \sigma_d^2)$ evaluated at the point z . Note that the Normal is quite flexible when used in this context. The Normal can (by choosing a very large variance) approximate closely the $d(z_s^*, z_{s-1}^*) = z_s - z_{s-1}$ distance function which is simply the distance between adjacent observations of the index variable. But it also can (by setting μ_d to τ and choosing a very small value for σ_d^2) approximate closely a structural break at time τ (or an abrupt change in the pass through coefficient when inflation expectations hit a threshold, τ). Intermediate values of the variance of the Normal would allow for a smooth change in the pass through coefficient around a threshold determined by μ_d . The advantage of our approach is that we treat μ_d, σ_d^2 as unknown parameters and, thus, the precise shape of the distance function is estimated from the data.

The interested reader is referred to Koop and Potter (2007) for a great deal of additional justification for this flexible parametric model. Suffice it to note here that it can uncover any plausible relationship between the

pass through coefficient and short term inflation expectations or any time variation in the coefficient. But it does so in a way where the data decides what the relationship is rather than imposing a particular relationship (as is done in standard regression-based approaches).

The previous material sketches out the basic concepts involved in our modeling strategy. To complete our empirical specification, we must discuss our treatment of the conditional variance. So far we have said nothing about stochastic volatility even though this is potentially important in studies of inflation dynamics (see, e.g., Stock and Watson, 2007). We give the conditional variance the same flexible parametric treatment that we give to the conditional mean. This will nest a standard stochastic volatility model, but be much more flexible. We also add an intercept to our model. Thus, our complete model is given by:

$$\Delta y_s = (1, \Delta x_s) \beta_s + \varepsilon_s, \quad (5)$$

where β_s is now a 2×1 vector containing an intercept and the pass through coefficient. Furthermore,

$$\beta_s = \beta_{s-1} + v_s \quad (6)$$

and $v_s \sim N \left[0, \sqrt{d_1(z_s, z_{s-1})} Q \right]$. The error in the measurement equation, (5). has the form,

$$\varepsilon_r = \xi_r \exp \left(\frac{1}{2} \alpha_r \right) \quad (7)$$

where $\xi_r \sim N(0, 1)$,

$$\alpha_r = \alpha_{r-1} + \eta_r \quad (8)$$

and $\eta_r \sim N \left[0, \sqrt{d_2(z_r, z_{r-1})} \sigma_\eta^2 \right]$. The errors, ξ_t , v_t and η_t , are independent at all leads and lags and are independent of one another. Note that the previous equations involve both s and r subscripts. These denote two different re-orderings in time according to two potentially different index variables. So, for instance, it is possible that the pass through coefficient varies with the level of expected inflation, but the conditional variance varies with time (i.e. has a standard stochastic volatility specification). The set of index variables is listed above. Finally, note that we allow for different distance functions for the conditional mean and variance.

Further details of our econometric methods are given in the Technical Appendix.

4 Empirical Results

This section presents empirical results using daily US data from January 2, 2003 through June 9, 2008. The short term (2-5 year) and long term (9-10 year) inflation expectations variables were calculated from the US real and nominal Treasury security yields (from the TIPS market) as described in Section 2. Potter and Rosenberg (2007) provide more detail of the calculations and a discussion of the TIPS market. The choices of starting date and 2-5 (9-10) as our definitions of short term (long term) are motivated by our desire to use data from highly liquid markets.

Our econometric models are as described in the previous section. Remember that we consider five different orderings (i.e. five different choices for the index variable) and different orderings can apply to the parameters characterizing the conditional mean and conditional variance. Accordingly, we have 25 different variants of our model. We note in passing that we experimented with restricted versions of our model where the conditional variance is constant (but nonlinearities can occur in the conditional mean) and the standard homoskedastic linear model but these were strongly rejected by the data and, hence, we do not present results for them here.

Note also that our approach allows for model selection (i.e. choosing the single model most supported by the data) or model averaging (i.e. averaging across all models with weights proportional to the marginal likelihoods). We present results using both approaches.

We present results for a subjectively-elicited but relatively noninformative prior. Furthermore, in the Empirical Appendix we present results for two training sample priors. One uses the first 200 observations to calibrate the prior while the other uses the last 200 observations. Full details on our priors are provided in the Technical Appendix.

4.1 Model Comparison

We begin in Table 1 with evidence on which models are supported by the data.⁴ With regards to the error variance (i.e. α_t) a strong story is obtained. Models which allow for stochastic volatility of a conventional sort (i.e. with data in natural time ordering) clearly receive much more support than any other models (i.e. models with other nonlinear patterns in their conditional variance). But there is some uncertainty over which nonlinear pattern in the conditional mean fits the data better. The model with the highest marginal likelihood has $z_t = \left| e^{i_t^{(2,5)}} - \bar{\pi} \right|$. Hence, if we select a single model we would choose the lagged absolute deviation from $\bar{\pi}$ as our index variable. But there is nearly as much support for some of the other index variables. Indeed, when doing Bayesian model averaging all of the index definitions for ordering β_t will receive non-negligible weight (but for α_t , the natural time ordering will receive virtually 100% of the weight).

⁴For the non-Bayesian reader, note the logs of marginal likelihoods can be interpreted in a similar fashion to information criteria. Indeed, the Schwarz criterion is an asymptotic approximation to the log marginal likelihood.

Table 1: Marginal Likelihoods for Different Orderings		
Index Variable for α_t	Index Variable for β_t	Log of Marginal Likelihood
t	t	2445.7
t	$ei_t^{(2,5)}$	2448.7
t	$ ei_t^{(2,5)} - \bar{\pi} $	2446.6
t	$ei_{t-1}^{(2,5)}$	2449.5
t	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2449.7
$ei_t^{(2,5)}$	t	2413.9
$ei_t^{(2,5)}$	$ei_t^{(2,5)}$	2407.0
$ei_t^{(2,5)}$	$ ei_t^{(2,5)} - \bar{\pi} $	2418.0
$ei_t^{(2,5)}$	$ei_{t-1}^{(2,5)}$	2401.0
$ei_t^{(2,5)}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2411.4
$ei_t^{(2,5)} - \bar{\pi}$	t	2421.5
$ei_t^{(2,5)} - \bar{\pi}$	$ei_t^{(2,5)}$	2409.1
$ei_t^{(2,5)} - \bar{\pi}$	$ ei_t^{(2,5)} - \bar{\pi} $	2418.3
$ei_t^{(2,5)} - \bar{\pi}$	$ei_{t-1}^{(2,5)}$	2405.5
$ei_t^{(2,5)} - \bar{\pi}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2418.2
$ei_{t-1}^{(2,5)}$	t	2426.1
$ei_{t-1}^{(2,5)}$	$ei_t^{(2,5)}$	2419.7
$ei_{t-1}^{(2,5)}$	$ ei_t^{(2,5)} - \bar{\pi} $	2414.4
$ei_{t-1}^{(2,5)}$	$ei_{t-1}^{(2,5)}$	2424.2
$ei_{t-1}^{(2,5)}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2422.4
$ei_{t-1}^{(2,5)} - \bar{\pi}$	t	2420.5
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ei_t^{(2,5)}$	2421.4
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ ei_t^{(2,5)} - \bar{\pi} $	2425.8
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ei_{t-1}^{(2,5)}$	2419.0
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2426.7

4.2 Features of Interest

In this section, we present posterior properties of some features of interest. Given that the main macroeconomic issues hinge on the pass-through coefficient, β_t , most of the empirical results in this section relate to them.

Figure 1 summarizes the posteriors of the pass through coefficients for a benchmark case (the standard TVP model) along with the two models using the index variables which yielded the highest log marginal likelihoods. These are the cases where $z_t = ei_{t-1}^{(2,5)}$ and $z_t = |ei_{t-1}^{(2,5)} - \bar{\pi}|$. For the conditional variance, we always have $z_t = t$ (i.e. conventional stochastic volatility) since this receives overwhelming support from the data. Note that the three panels of Figure 1 plot the inflation pass through coefficient against the appropriate index variable so that each panel has a different X-axis. The panels contain a point estimate (the posterior mean) along with credible interval (the interval between the 16th and 84th percentiles of the posterior).

Note that the inflation pass through coefficient is small everywhere. For no value of any index variable does its point estimate go above 0.3 and the credible intervals are fairly tight around the point estimate. Thus, there is strong evidence against the unmoored inflation expectations hypothesis. However, the inflation pass through coefficient is also not constant and near zero as would be suggested by firmly anchored inflation expectations.

Our results are most consistent with contained inflation expectations. Remember that this hypothesis involved the existence of an interval, inside which investors expected the Fed to keep inflation. It implied that the inflation pass through coefficient should decrease near the bounds of the interval. However, near the middle of the interval, the coefficient would be quite large (potentially even being near one). The results for our most probable model (i.e. the one where $|ei_{t-1}^{(2,5)} - \bar{\pi}|$ is the index variable) are consistent with this pattern (see the bottom panel of Figure 1). When short term inflation expectations are near $\bar{\pi}$ there is a fair degree of persistence (i.e. the inflation pass through coefficient is around 0.2), but if they deviate more than one percent from $\bar{\pi}$ then the persistence vanishes (i.e. the pass through coefficient is virtually zero). This pattern is also consistent with a less rigid form of the contained inflation expectations hypothesis where investors believe that the Fed will not act strongly to correct small deviations from $\bar{\pi}$ (e.g. less than 1% deviations), but will increasingly tend to act as deviations become larger. In summary, if we select a single model with the highest marginal likelihood,

we are not finding evidence for either unmoored nor firmly anchored inflation expectations, but rather something in between where inflation expectations are contained.

Figure 2 present point estimates of the inflation pass through coefficients, but averaged over all models (and plotted with the X-axis ordered in three different ways).⁵ Note that, unlike Figure 1, the three panels in Figure 2 do not correspond to different models, but instead are presenting the same information but ordered in different ways. In the bottom panel of Figure 2 we see again the pattern consistent with contained inflation expectations.

Figures 1 and 2 highlight two qualifications of this story. First, results using $z_t = ei_{t-1}^{(2,5)}$ as the index variable indicate that the pass through coefficient is near zero when short term inflation expectations are low (i.e. below $\bar{\pi}$), but is higher when they are high (above $\bar{\pi}$). We stress that this index variable yields a lower marginal likelihood than using $|ei_{t-1}^{(2,5)} - \bar{\pi}|$, so we place less weight on this model. Nevertheless, it suggests that the evidence of contained inflation expectations is mostly coming from periods where $ei_{t-1}^{(2,5)} < \bar{\pi}$. Looking at the raw data, $\bar{\pi}$ is approximately the mean of $ei_t^{(2,5)}$, the distribution of $ei_t^{(2,5)}$ about this mean is asymmetric with more extreme deviations occurring below than above the mean. It is these larger deviations below the mean (which often occur near the beginning of our sample) which are playing a large role in our finding of contained inflation expectations.

Second, the results using $z_t = t$ (a model which, of the three plotted in Figure 1, has much the lowest marginal likelihood) indicate that the inflation pass through coefficient was lower at the beginning of our sample period. This is consistent with the final statement in the preceding paragraph.

These results are obtained using a subjectively-elicited, but relatively noninformative prior. In order to investigate possible prior sensitivity, we have also produced results using two training sample priors. These are in the Empirical Appendix and the interested reader can look at them in detail. Here we note that the basic story outlined above is robust. That is, we are finding no evidence for either unmoored or firmly anchored inflation expectations. Inflation expectations do seem to be better characterized as being contained. However, observations at the beginning of the sample do play a

⁵The rather irregular patterns of the posterior means in some places are due to our averaging over models which order the data in very different ways. To keep the graphs uncluttered we do not include credible intervals (their approximate width is similar to those in Figure 1).

disproportionate role in this finding.

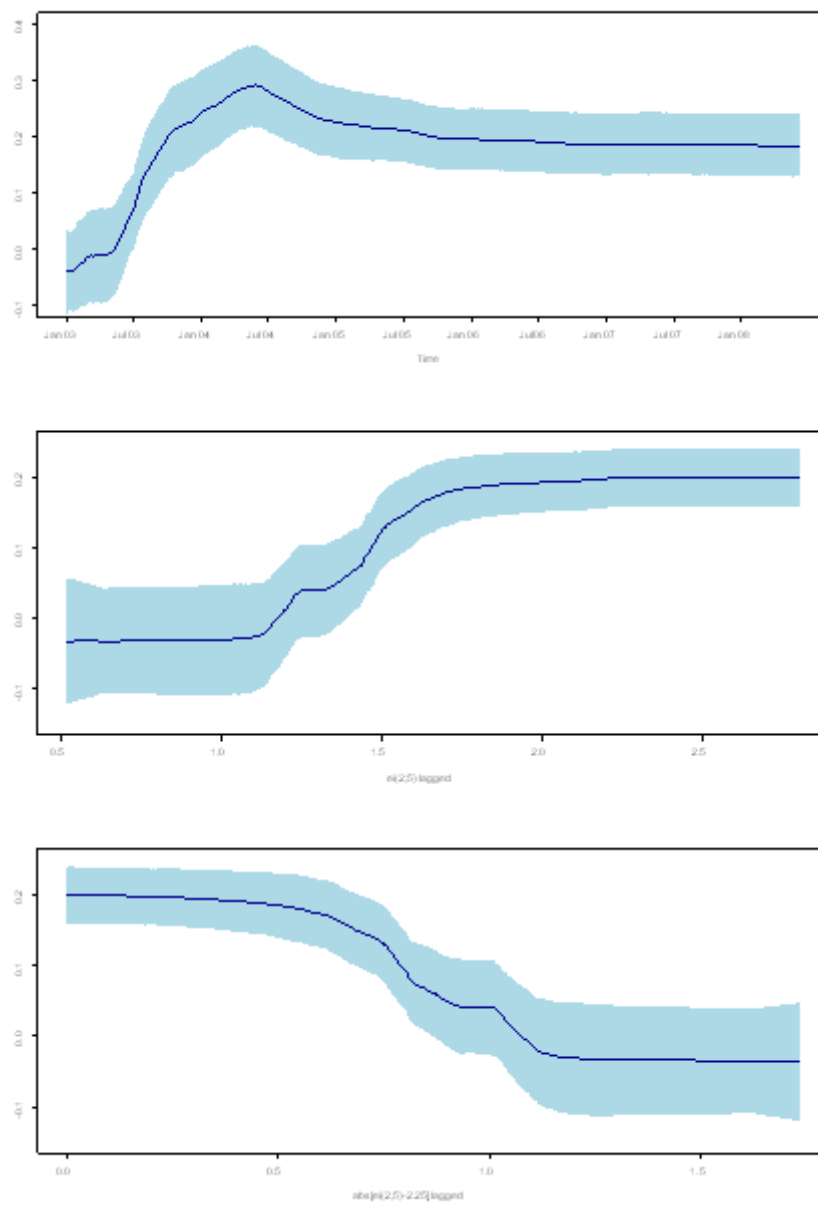


Figure 1: Posterior Means of Pass-through Coefficient for Models with Different Index Variables (with 16th-84th percentile bands)

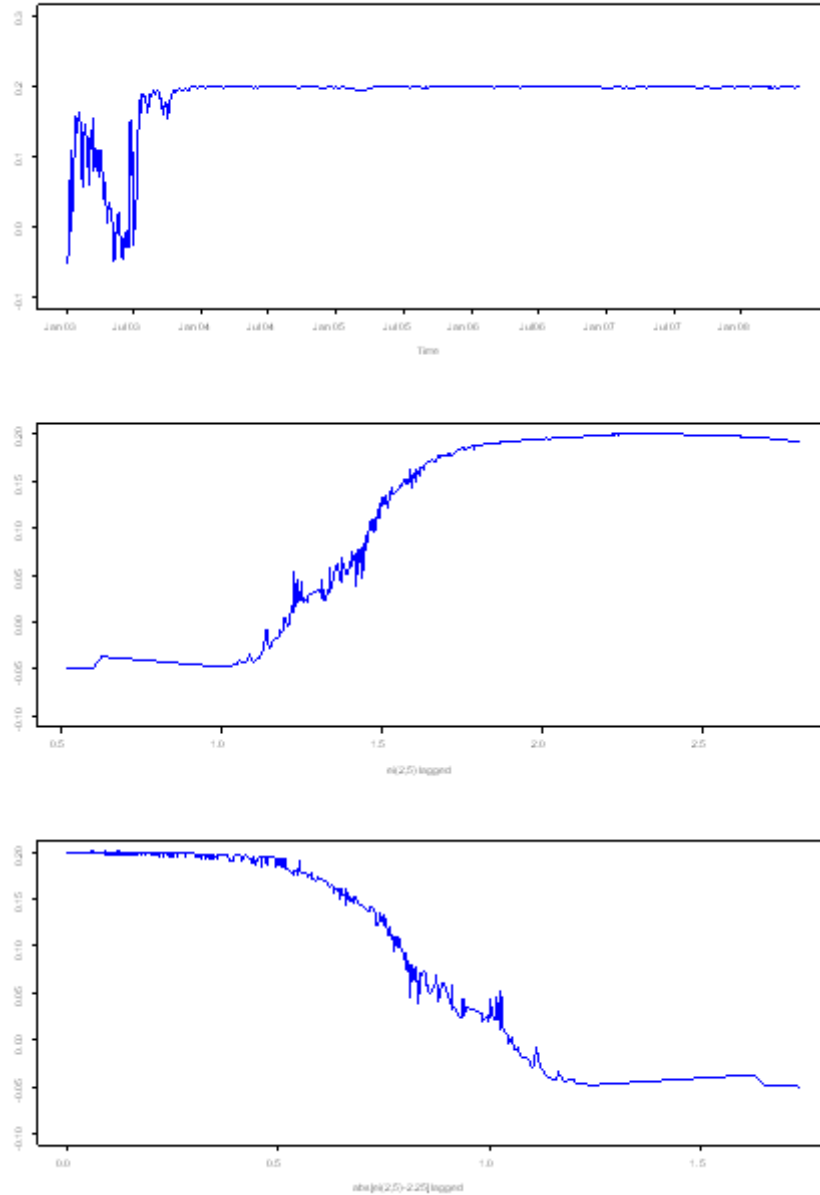


Figure 2: Posterior Means of Pass-through Coefficient. Averaged Over All Models. Plotted Against Different Index Variables

The main stories in our paper relate to the inflation pass through coefficient. However, our empirical results are clearly indicating that the error

variance is not constant over time. In fact, Table 1 shows that there is strong evidence of stochastic volatility of a conventional sort in this data set. Accordingly, we present Figure 3 which plots posterior results relating to the conditional variance in a similar format to Figure 1. That is, the three panels of Figure 3 correspond to the same three models as in Figure 1 (although all of these use the same index for the conditional variance and accordingly all have time as the X-axis). Remember that Table 1 indicates strong support for conventional stochastic volatility (so e.g. there is no support for the idea that the error variance depends on the level of short term inflation expectations). All three panels of Figure 3 look very similar and indicate that the error variance has been fluctuating substantially over time.

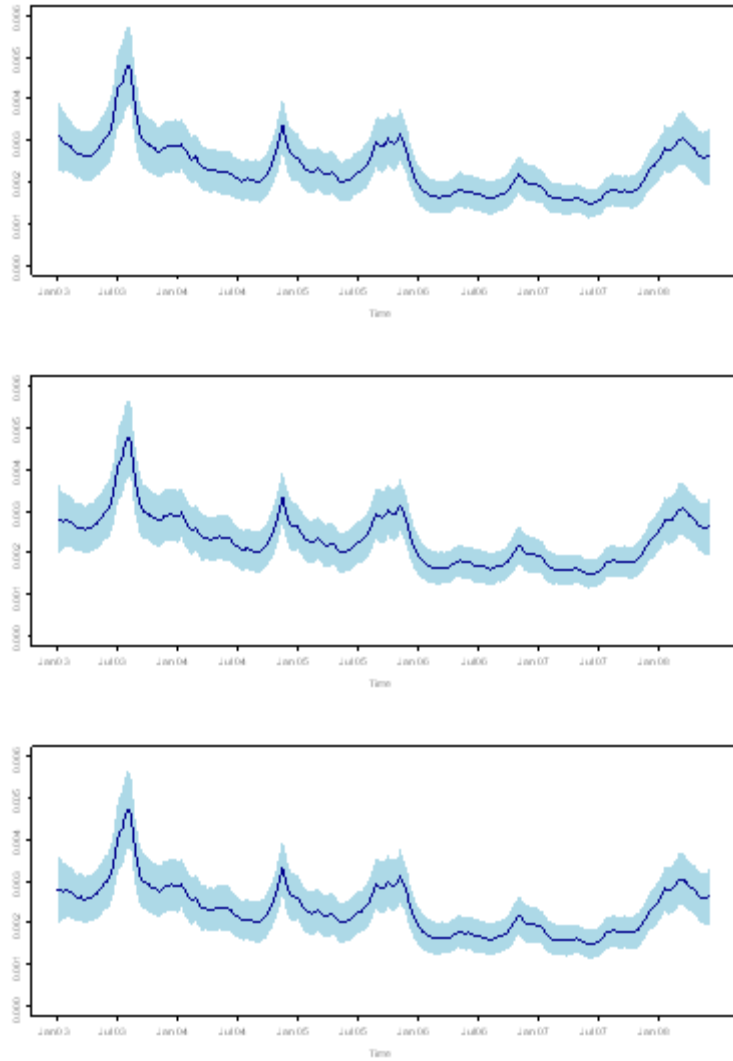


Figure 3: Posterior Means of Conditional Variances for Models with Different Index Variables (with 16th-84th percentile bands)

5 Conclusions

In this paper, we have investigated the relationship between short-term and long-term inflation compensation. We focus on the inflation pass-through coefficient which measures how changes in short-term expectations impact on long term expectations. We use a flexible parametric approach which allows this pass through coefficient to vary over time or vary with the level of short term inflation expectations or vary with absolute deviations of inflation expectations from a central value. Furthermore, the relationship between the pass through coefficient and any of these index variables can have virtually any functional form.

Our empirical results indicate that, although there is a fair degree of model uncertainty, that most support goes to a model where the inflation pass through coefficient depends on the absolute deviation of short term inflation expectations from a central value and that the shape of this relationship is consistent with a theoretical model of contained inflation expectations.

References

- Chib, S. and Greenberg, E. (1995). "Understanding the Metropolis-Hastings algorithm," *The American Statistician*, 49, 327-335.
- Durbin, J. and Koopman, S. (2002). "A simple and efficient simulation smoother for state space time series analysis," *Biometrika*, 89, 603-616.
- Faust, J. and Henderson, D. (2004). "Is inflation targetting best-practice monetary policy?" Federal Reserve Bank of St. Louis Review, 86, 117.
- Gelfand, A. and Dey, D. (1994). "Bayesian model choice: Asymptotics and exact calculations," *Journal of the Royal Statistical Society Series B*, 56, 501-514.
- Gurkaynak, R., Levin, A. and Swanson, E. (2006). "Does inflation targetting anchor long-run inflation expectations?" Federal Reserve Bank of San Francisco, working paper 2006-09.
- Harvey, A. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press: Cambridge.
- Harvey, A. and Koopman, S. (2000). "Signal extraction and the formulation of unobserved components models," *Econometrics Journal*, 3, 84-107.
- Kim, S., Shephard, N. and Chib, S. (1998). "Stochastic volatility: Likelihood inference and comparison with ARCH models," *Review of Economic Studies*, 65, 361-393.
- Koop, G. (2003). *Bayesian Econometrics*, John Wiley: Chichester.
- Koop, G. and Potter, S. (2007). "A flexible approach to parametric inference in nonlinear and time varying time series models," manuscript available at <http://personal.strath.ac.uk/gary.koop/>.
- Potter, S. (2007). "A bounded model of inflation," Federal Reserve Bank of New York, Staff Reports.
- Potter, S. and Rosenberg, J. (2007). "Are U.S. inflation expectations anchored, contained or unmoored?" manuscript.
- Stock, J. and Watson, M. (2007). "Why has U.S. inflation become harder to forecast?" *Journal of Money, Credit and Banking*, 39, 3-33.
- Tsay, R. (1989). "Testing and modeling threshold autoregressive processes," *Journal of the American Statistical Association*, 84, 231-240

Technical Appendix

The Markov chain Monte Carlo (MCMC) algorithm for the flexible parametric model is fully described in Koop and Potter (2007) and the reader is referred there for complete details. Here we just sketch out the basic ideas. The advantage of our approach is that we can draw upon standard state space algorithms. We begin by describing how to do MCMC for a given choice of two index variables (one controlling variation in the conditional mean, the other in the conditional variance).

The model is given in (5), (6), (7) and (8). To draw β_t for $t = 1, \dots, T$ (or when we re-order the data, β_s for $s = 1, \dots, T$), we use the algorithm of Durbin and Koopman (2002). We use the algorithm of Kim, Shephard and Chib (1998) (when necessary using suitably re-ordered data) to draw α_r for $t = 1, \dots, T$ (or α_r for $r = 1, \dots, T$ when we re-order the data).

For Q and σ_η^{-2} we use inverted-Wishart and Gamma priors, respectively. Note that, conditional on the states, the state equations reduce to simplified variants of linear regression models. Thus, for our priors and conditional on draws of the states, the posterior for Q takes the usual inverted-Wishart form (see, e.g., Koop, 2003, pages 140-141) and the posterior of σ_η^{-2} takes the usual Gamma form (see, e.g., Koop, 2003, pages 61-62).

For any given ordering, we need to draw the parameters of the distance functions. The number of parameters in these is only four. That is, we have a μ_d and a σ_d^2 for the distance function in (6) as well as (8). Random Walk Chain Metropolis-Hastings algorithms (see, e.g., Chib and Greenberg, 1995) work well in low-dimensional models such as this and we use such an algorithm. Koop and Potter (2007) gives the formulae for the acceptance probabilities.

The preceding material described our MCMC algorithm, given a particular choice of index variables. To average over index variables (or select a particular choice), we must calculate the marginal likelihood for each choice. We use the approach of Gelfand and Dey (1994) to do so.

With regards to the prior, we use a subjective (but relatively noninformative) prior as well as two training sample priors. The subjective prior is similar to that motivated in Koop and Potter (2007) and the reader is referred to that paper for a detailed justification. Note that all of our models have the same dimension, so some problems with the use of relatively noninformative priors (e.g. Bartlett's paradox) do not arise. The key priors are those for the initial conditions, the error variances in the state equations and those for the parameters of the distance functions. For these, we use the

same priors for every choice of index variables.

In particular, we assume priors of the form:⁶

$$Q \sim iW(\underline{Q}, \underline{q}),$$

$$\sigma_\eta^{-2} \sim G(\underline{a}_\eta, \underline{b}_\eta),$$

$$\alpha_1 \sim N(\underline{\alpha}, \underline{V}_\alpha),$$

$$\beta_1 \sim N(\underline{\beta}, \underline{V}_\beta),$$

$$\mu_d \sim N(\underline{\mu}, \underline{V}_\mu)$$

and

$$\sigma_d^{-2} \sim G(\underline{a}_d, \underline{b}_d)$$

The subjective prior is obtained setting $\underline{\alpha} = -5$, $\underline{V}_\alpha = 10$, $\underline{\beta} = 0$ and $\underline{V}_\beta = 10 \cdot I_2$. Furthermore, $\underline{a}_\eta = 400$, $\underline{b}_\eta = 2$, $\underline{q} = 50$ and $\underline{Q} = \begin{bmatrix} 0.001 & 0.000 \\ 0.000 & 0.005 \end{bmatrix}$.

For the training sample priors, we use either the first or the last 200 observations. Based on OLS results using either training sample we specify the prior parameters for the initial conditions as $\underline{\alpha} = \log(\hat{\sigma}_{OLS}^2)$, $\underline{V}_\alpha = 10$, $\underline{\beta} = \hat{\beta}_{OLS}$ and $\underline{V}_\beta = 10 \cdot \text{var}(\hat{\beta}_{OLS})$. We then set $\underline{a}_\eta = 400$, $\underline{b}_\eta = 2$, $\underline{q} = 50$ and $\underline{Q} = 0.01(50 - 3) \text{var}(\hat{\beta}_{OLS})$.

The hyperparameters of both distance function priors are the same for the training sample priors and the subjective prior and set at $\underline{\mu} = 0$, $\underline{V}_\mu = 1$, $\underline{a}_d = 1.0$ and $\underline{b}_d = 1.0$.

⁶The Gamma is parameterized so that $G(a, b)$ has mean $\frac{a}{b}$ and variance $\frac{a}{b^2}$. The inverted Wishart is parameterized so that $iW(A, a)$ has expected value $\frac{A}{a - p - 1}$ where p is the dimension.

Empirical Appendix

In this appendix, we present posterior results for our two training sample priors. Figures A.1 and A.2 are in the same format as Figures 1 and use the last 200 observations and the first 200 observations, respectively, to calibrate the prior. Table A.1 is comparable to Table 1 and contains log marginal likelihoods for different index variables.

Consider first the training sample using the last 200 observations. Table A.1 indicates that the most preferred model has the index variable being the absolute deviation from $\bar{\pi}$. Figure A.1 looks very similar to Figure 1. This strongly supports the idea that our results are robust to prior choice.

Results using the first 200 observations, though, exhibit less robustness. The models with highest marginal likelihoods are not those where the index is the absolute deviation from $\bar{\pi}$, but rather simply the level of short term inflation expectations (or its lag). This is consistent with the idea discussed in the body of the paper: That the first observations (which are set aside for this training sample, but are included with the other priors) are important in providing support for the model where the index is the absolute deviation from $\bar{\pi}$. Nevertheless, most aspects of our basic story (i.e. that inflation expectations seem not to be unmoored nor anchored, but rather are contained in some way) are left unchanged.

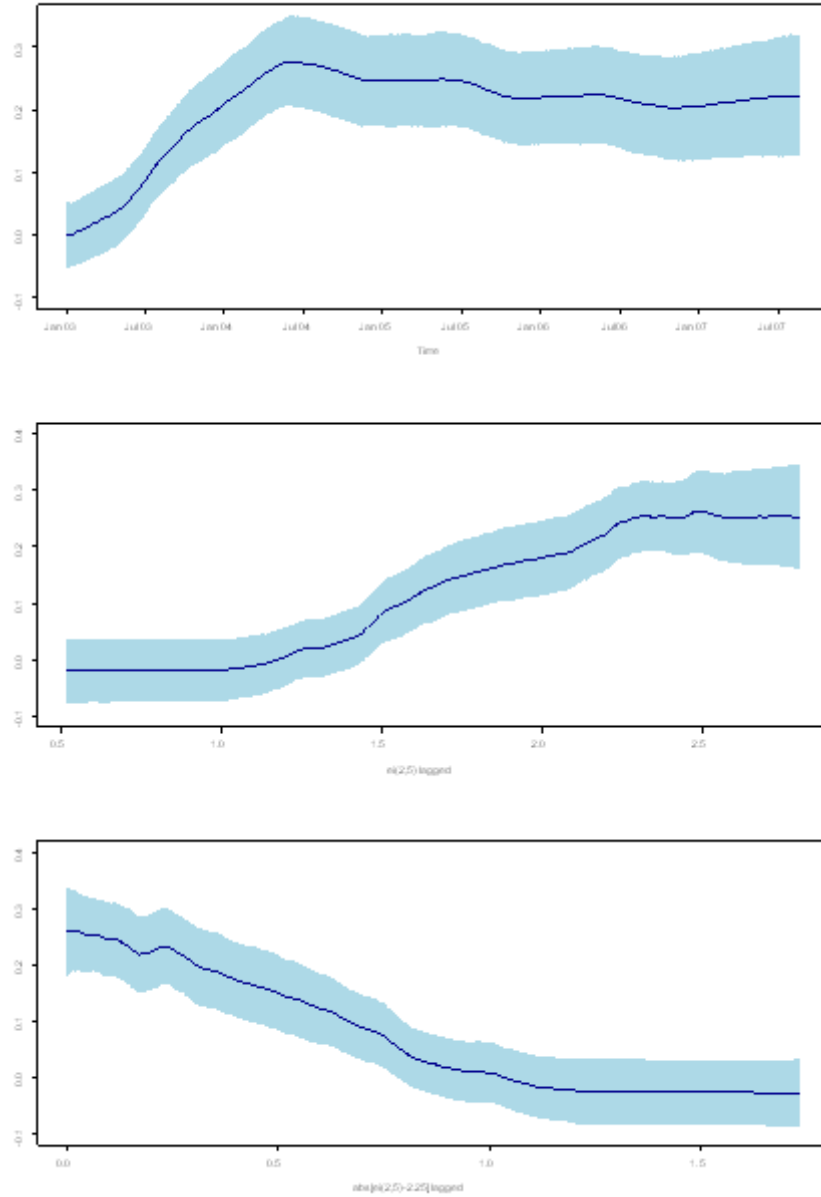


Figure A.1: Posterior of Pass-through Coefficient for Models with Different Index Variables (Training sample last 200 observations)

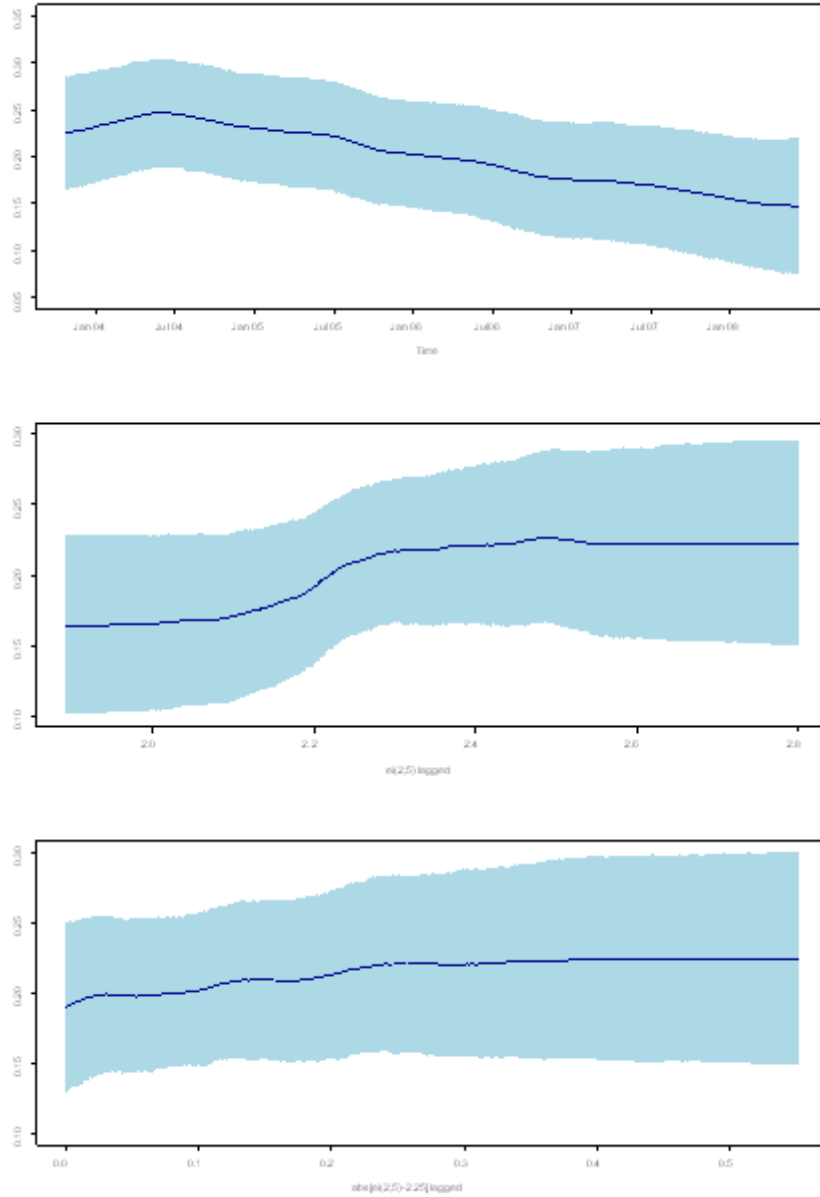


Figure A.2: Posterior of Pass-through Coefficient for Models with Different Index Variables (Training sample first 200 observations)

Table A.1: Marginal Likelihoods for Different Orderings

		Training sample:	Last 200 obs.	First 200 obs.
Index Variable for α_s	Index Variable for β_s	Log of Marginal Likelihood	Log of Marginal Likelihood	Log of Marginal Likelihood
t	t	2122.3	2164.3	2164.3
t	$ei_t^{(2,5)}$	2133.9	2173.6	2173.6
t	$ ei_t^{(2,5)} - \bar{\pi} $	2138.4	2166.2	2166.2
t	$ei_{t-1}^{(2,5)}$	2136.0	2175.7	2175.7
t	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2136.2	2169.6	2169.6
$ei_t^{(2,5)}$	t	2108.3	2150.3	2150.3
$ei_t^{(2,5)}$	$ei_t^{(2,5)}$	2107.2	2142.3	2142.3
$ei_t^{(2,5)}$	$ ei_t^{(2,5)} - \bar{\pi} $	2107.7	2144.5	2144.5
$ei_t^{(2,5)}$	$ei_{t-1}^{(2,5)}$	2107.4	2147.8	2147.8
$ei_t^{(2,5)}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2109.9	2145.9	2145.9
$ei_t^{(2,5)} - \bar{\pi}$	t	2098.5	2146.6	2146.6
$ei_t^{(2,5)} - \bar{\pi}$	$ei_t^{(2,5)}$	2099.5	2154.0	2154.0
$ei_t^{(2,5)} - \bar{\pi}$	$ ei_t^{(2,5)} - \bar{\pi} $	2091.7	2143.3	2143.3
$ei_t^{(2,5)} - \bar{\pi}$	$ei_{t-1}^{(2,5)}$	2092.6	2151.8	2151.8
$ei_t^{(2,5)} - \bar{\pi}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2100.9	2150.2	2150.2
$ei_{t-1}^{(2,5)}$	t	2118.1	2150.8	2150.8
$ei_{t-1}^{(2,5)}$	$ei_t^{(2,5)}$	2114.6	2146.1	2146.1
$ei_{t-1}^{(2,5)}$	$ ei_t^{(2,5)} - \bar{\pi} $	2114.4	2142.8	2142.8
$ei_{t-1}^{(2,5)}$	$ei_{t-1}^{(2,5)}$	2117.2	2147.4	2147.4
$ei_{t-1}^{(2,5)}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2113.1	2148.7	2148.7
$ei_{t-1}^{(2,5)} - \bar{\pi}$	t	2101.5	2147.0	2147.0
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ei_t^{(2,5)}$	2097.3	2158.7	2158.7
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ ei_t^{(2,5)} - \bar{\pi} $	2100.9	2151.8	2151.8
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ei_{t-1}^{(2,5)}$	2091.6	2147.8	2147.8
$ei_{t-1}^{(2,5)} - \bar{\pi}$	$ ei_{t-1}^{(2,5)} - \bar{\pi} $	2098.8	2145.7	2145.7